arXiv:1610.07315v1 [math.NA] 24 Oct 2016

Discrete least-squares approximations over optimized downward closed polynomial spaces in arbitrary dimension *

Albert Cohen† Giovanni Migliorati‡ Fabio Nobile§

October 25, 2016

Abstract

We analyze the accuracy of the discrete least-squares approximation of a function u in multivariate polynomial spaces $\mathbb{P}_\Lambda := \text{span}\{y \mapsto y^\nu : \nu \in \Lambda\}$ with $\Lambda \subset \mathbb{N}_0^d$ over the domain $\Gamma := [-1,1]^d$, based on the sampling of this function at points $y^1, \dots, y^m \in \Gamma$. The samples are independently drawn according to a given probability density ρ belonging to the class of multivariate beta densities, which includes the uniform and Chebyshev densities as particular cases. Motivated by recent results on high-dimensional parametric and stochastic PDEs, we restrict our attention to polynomial spaces associated with *downward closed* sets Λ of *prescribed* cardinality n, and we optimize the choice of the space for the given sample. This implies, in particular, that the selected polynomial space depends on the sample. We are interested in comparing the error of this least-squares approximation measured in $L^2(\Gamma, d\rho)$ with the best achievable polynomial approximation error when using downward closed sets of cardinality n. We establish conditions between the dimension n and the size m of the sample, under which these two errors are proven to be comparable. Our main finding is that the dimension d enters only moderately in the resulting trade-off between m and n, in terms of a logarithmic factor $\ln(d)$, and is even absent when the optimization is restricted to a relevant subclass of downward closed sets, named *anchored* sets. In principle, this allows one to use these methods in arbitrarily high or even infinite dimension. Our analysis builds upon [2] which considered fixed and nonoptimized downward closed multi-index sets. Potential applications of the proposed results are found in the development and analysis of efficient numerical methods for computing the solution to high-dimensional parametric or stochastic PDEs, albeit not limited to this area.

Keywords: multivariate polynomial approximation, discrete least squares, convergence rate, best n-term approximation, downward closed set.

MSC: 41A10, 41A25, 41A50, 41A63, 65M70.

1 Introduction

In recent years it has become clear that many interesting engineering applications feature an intrinsic dependence on a large number of parameters $y_1, \dots, y_d$, leading to a major concentration of efforts in

*This research was supported by the Institut Universitaire de France, the ERC AdG BREAD, and the center for advanced modeling science (CADMOS).

†Sorbonne Universités, UPMC Univ Paris 06, CNRS, UMR 7598, Laboratoire Jacques-Louis Lions, 4, place Jussieu, Paris 75005, France. email: cohen@ann.jussieu.fr

‡Sorbonne Universités, UPMC Univ Paris 06, CNRS, UMR 7598, Laboratoire Jacques-Louis Lions, 4, place Jussieu, Paris 75005, France. email: migliorati@ann.jussieu.fr

§MATHICSE-CSQI, École Polytechnique Fédérale de Lausanne, Lausanne CH-1015, Switzerland. email: fabio.nobile@epfl.ch

the development and analysis of high-dimensional approximation methods. In many relevant situations, the parameters y_j are independent random variables distributed on intervals I_j according to probability measures $d\rho_j$. We are then typically interested in approximating a function

$$y = (y_1, \dots, y_d) \mapsto u(y),$$

depending on these parameters and measuring the error in $L^2(\Gamma, d\rho)$, where $\Gamma = I_1 \times \cdots \times I_d$ and $d\rho = d\rho_1 \otimes \cdots \otimes d\rho_d$. Up to a renormalization, we may assume that $I_j = [-1, 1]$ for all j, so that $\Gamma = [-1, 1]^d$. In certain situations, the number of parameters may even be countably infinite, in which case $\Gamma = [-1, 1]^{\mathbb{N}}$. Examples where such problems occur are recurrent in the numerical treatment of parametric and stochastic PDEs, where a fast and accurate approximation of the parameter-to-solution map over high-dimensional parameter sets is useful to tackle more complex optimization, control and inverse problems.

In such a context, the potential of specific high-dimensional *polynomial* approximation methods has been demonstrated in [10, 8, 20, 17, 2]. In these methods, the approximation is picked from a multivariate polynomial space

$$\mathbb{P}_\Lambda := \operatorname{span}\{y \mapsto y^\nu \,:\, \nu \in \Lambda\},$$

where Λ is a given finite subset of $\mathbb{N}_0^d$. In the case of countably many parameters, $d = \infty$, we replace $\mathbb{N}_0^d$ by the set of finitely supported sequences of nonnegative integers.

The set Λ is said to be *downward closed* if and only if

$$\nu \in \Lambda \quad \text{and} \quad \mu \le \nu \implies \mu \in \Lambda, \tag{1}$$

where $\mu \le \nu$ is meant component-wise as $\mu_i \le \nu_i$ for all i. Polynomial spaces $\mathbb{P}_\Lambda$ associated to downward closed index sets Λ have been studied in various contexts, see [3, 13, 12, 15, 16].

There exist two main approaches to polynomial approximation of a given function u based on pointwise evaluations. The first approach relies on *interpolation* of the function u at a given set of points $\{y^1, \dots, y^n\}$ where $n := \#(\Lambda) = \dim(\mathbb{P}_\Lambda)$, that is, find $v \in \mathbb{P}_\Lambda$ such that $v(y^i) = u(y^i)$ for $i = 1, \dots, n$. The second approach relies on *projection*, which aims at minimizing the $L^2(\Gamma, d\rho)$ error between u and its approximation in $\mathbb{P}_\Lambda$. Since the exact projection is not available, one typical strategy consists in using the discrete least-squares method, that is, solving the problem

$$\min_{v \in \mathbb{P}_\Lambda} \sum_{i=1}^m |v(y^i) - u(y^i)|^2,$$

where now $m > n$. Discrete least-squares methods are often preferred to interpolation methods when the observed evaluations are polluted by noise. Their convergence analysis has been studied in the general context of learning theory, see for example [11, 23, 14, 26, 27].

In recent years, an analysis of discrete least-squares methods has been proposed [2, 21, 17, 22], specifically targeted to the above described case of multivariate polynomial spaces associated with downward closed sets, in the case where the $d\rho_j$ are identical Jacobi-type measures. This analysis, which builds upon the general results from [7], gives conditions ensuring that, in the absence of noise in the pointwise evaluation of u, the accuracy of the discrete least-squares approximation is comparable to the best approximation error achievable in $\mathbb{P}_\Lambda$, that is,

$$e_\Lambda(u) := \inf_{v \in \mathbb{P}_\Lambda} \|u - v\|_{L^2(\Gamma, d\rho)}.$$

These conditions are stated in terms of a relation between the size m of the sample and the dimension n of $\mathbb{P}_\Lambda$. A similar analysis also covers the case of an additive noise in the evaluation of the samples, which results in additional terms in the error estimate, see *e.g.* [22].

One remarkable result from the above analysis is that the conditions ensuring that the least-squares method has accuracy comparable to $e_\Lambda(u)$ only involve the dimension of $\mathbb{P}_\Lambda$. These conditions are independent of the specific shape of the set Λ (as long as it is downward closed), and in particular independent of the dimension d.

The possibility of using arbitrary sets Λ is critical in the context of parametric PDEs in view of the recent results on high-dimensional polynomial approximation obtained in [10, 4, 8]. These results show that for relevant classes of parametric PDEs, the functions $y \mapsto u(y)$ describing either the full solution or scalar quantities of interest can be approximated with convergence rates $\mathcal{O}(n^{-s})$ which are *independent* of the parametric dimension d, when using polynomial spaces $\mathbb{P}_{\Lambda_n}$ associated to *specific* sequences of downward closed multi-index sets $(\Lambda_n)_{n\geq 1}$ with $\#(\Lambda_n) = n$. In summary, we have

$$e_n(u) := \min_{\#(\Lambda)=n} e_\Lambda(u) \leq Cn^{-s}, \tag{2}$$

where the minimum is taken over all downward closed sets of given cardinality n.

For each value of n, the optimal set Λ_n is the one that achieves the minimum in (2) among all downward closed Λ of cardinality n. This set is unknown to us when observing only the samples $u(y^i)$. Therefore, a legitimate objective is to develop least-squares methods for which the accuracy is comparable to the quantity $e_n(u)$.

In this paper, we discuss least-squares approximations on multivariate polynomial spaces for which the choice of Λ is optimized based on the sample. In particular we prove that the performance of such approximations is comparable to the quantity in (2), under a relation between m and n where the dimension d enters as a logarithmic factor. Furthermore, we show that this logarithmic dependence on d can be fully removed by considering a more restricted class of downward closed sets called *anchored sets*, for which similar approximation rates as in (2) can be achieved. The resulting least-squares methods are thus immune to the curse of dimensionality.

The outline of the paper is the following: in Section 2 we introduce the notation and briefly review some of the previous results achieved in the analysis of discrete least squares on *fixed* multivariate polynomial spaces. In Section 3 we present the main results of the paper concerning discrete least-squares approximations on *optimized* polynomial spaces. Our analysis is based on establishing upper bounds on the number of downward closed or anchored sets of a given cardinality, or on the cardinality of their union.

The selection of the optimal polynomial space is based on minimizing the least-squares error among all possible choices of downward closed or anchored sets of a given cardinality n. Let us stress that in the form of an exhaustive search, this task becomes computationally intensive when n and d are simultaneously large. Our results should therefore mainly be viewed as a benchmark in arbitrary dimension d for assessing the performance of fast selection algorithms, such as greedy algorithms, that still need to be developed and analyzed in this context. A general discussion on alternate selection strategies with reasonable computational cost is presented in the final section §4.

2 Discrete least-squares approximation by multivariate polynomials

In this section we introduce some useful notation, and recall from [2] the main results achieved for the analysis of the stability and accuracy of discrete least-squares approximations in multivariate polynomial spaces.

2.1 Notation

In any given dimension $d \in \mathbb{N}$, we consider the domain $\Gamma := [-1,1]^d$, and for some given real numbers $\theta_1, \theta_2 > -1$, the tensorized Jacobi measure

$$d\rho = \otimes_{j=1}^d d\beta,$$

with density

$$\rho(y) := \prod_{j=1}^d \beta(y_j),$$

where

$$d\beta = \beta(t)dt := c(1-t)^{\theta_1}(1+t)^{\theta_2}dt, \quad c := \left(\int_{-1}^1 (1-t)^{\theta_1}(1+t)^{\theta_2}dt\right)^{-1}.$$

We may also consider the case $\Gamma := [-1,1]^{\mathbb{N}}$ for which $d = +\infty$ and $d\rho$ is the Jacobi measure defined over Γ in the usual manner. We denote by $L^2(\Gamma, d\rho)$ the Hilbert space of real-valued square-integrable functions with respect to $d\rho$ and denote by $\|\cdot\|$ the associated norm, i.e.

$$\|v\| := \left(\int_\Gamma |v(y)|^2 d\rho(y)\right)^{1/2}.$$

Moreover, let $\mathcal{F}$ be defined as the set $\mathbb{N}_0^d$, where $\mathbb{N}_0 := \{0, 1, 2, \dots\}$, in the case $d < +\infty$, or as the countable set of all finitely supported sequences from $\mathbb{N}_0^{\mathbb{N}}$ in the case $d = +\infty$. For any $\nu \in \mathcal{F}$ we define

$$\mathrm{supp}(\nu) := \{j \geq 1 : \nu_j \neq 0\},$$

and for any multi-index set $\Lambda \subseteq \mathcal{F}$ we define

$$\mathrm{supp}(\Lambda) := \cup_{\nu \in \Lambda} \mathrm{supp}\{\nu\}.$$

We say that a coordinate y_j is *active* in the space $\mathbb{P}_\Lambda$ when $j \in \mathrm{supp}(\Lambda)$.

For the given real parameters $\theta_1, \theta_2 > -1$, we introduce the family $(J_n)_{n \geq 0}$ of univariate orthonormal Jacobi polynomials associated with the measure $d\beta$, and their tensorized counterpart

$$J_\nu(y) = \prod_{j \geq 1} J_{\nu_j}(y_j), \quad y = (y_j)_{j \geq 1},$$

for any $\nu \in \mathcal{F}$. The $(J_\nu)_{\nu \in \mathcal{F}}$ are an $L^2(\Gamma, d\rho)$-orthonormal basis. Particular instances of these polynomials are tensorized Legendre polynomials when $\theta_1 = \theta_2 = 0$ and tensorized Chebyshev polynomials of the first kind when $\theta_1 = \theta_2 = -1/2$.

In the present paper we focus on finite multi-index sets Λ which are downward closed in the sense of (1). We also say that a polynomial space $\mathbb{P}_\Lambda$ is downward closed when it is associated with a downward closed multi-index set $\Lambda \subset \mathcal{F}$. Recall that $\mathbb{P}_\Lambda$ has been defined as the span of the monomials $y \mapsto y^\nu$ for $\nu \in \Lambda$. Therefore it admits $(J_\nu)_{\nu \in \Lambda}$ as an $L^2(\Gamma, d\rho)$-orthonormal basis in the case of Λ downward closed. Sometimes we enumerate the indices ν using the lexicographical ordering, and denote this basis by $(\psi_k)_{k=1,\dots,n}$, where

$$n := \#(\Lambda) = \dim(\mathbb{P}_\Lambda).$$

Given a finite downward closed multi-index set $\Lambda \subset \mathcal{F}$, we would like to approximate the target function $u : \Gamma \to \mathbb{R}$ in the L^2 sense, using the noiseless evaluations $(u(y^i))_{i=1,\dots,m}$ of u at the points

$(y^i)_{i=1,\dots,m}$, where the y^i are i.i.d. random variables distributed according to ρ. We define the continuous L^2 projection of u on the polynomial space $\mathbb{P}_\Lambda$ as

$$\Pi_\Lambda u := \operatorname*{argmin}_{v\in\mathbb{P}_\Lambda} \|u-v\|,$$

and the discrete least-squares approximation $\Pi_\Lambda^m u$ as

$$\Pi_\Lambda^m u := \operatorname*{argmin}_{v\in\mathbb{P}_\Lambda} \|u-v\|_m, \tag{3}$$

where we have used the notation

$$\|v\|_m := \left(\frac{1}{m}\sum_{i=1}^m v(y^i)^2\right)^{\frac{1}{2}}.$$

We introduce the $m\times\#(\Lambda)$ design matrix $\mathbf{D}$ and the vector $\boldsymbol{b}\in\mathbb{R}^m$ whose entries are given by $\mathbf{D}_{i,k}=\psi_k(y^i)$ and $\boldsymbol{b}_i=u(y^i)$. We define the Gramian matrix $\mathbf{G}:=m^{-1}\mathbf{D}^{\mathrm{T}}\mathbf{D}$. The discrete least-squares projection in (3) is then given by

$$\Pi_\Lambda^m u = \sum_{k=1}^{\#(\Lambda)} \mathbf{a}_k\psi_k,$$

where the vector $\mathbf{a}=(\mathbf{a}_k)\in\mathbb{R}^{\#(\Lambda)}$ is the solution to the normal equations

$$\mathbf{G}\mathbf{a} = m^{-1}\mathbf{D}^{\mathrm{T}}\boldsymbol{b}.$$

2.2 Previous results on the stability and accuracy of discrete least squares

We introduce the quantity

$$K(\mathbb{P}_\Lambda) := \sup_{y\in\Gamma}\sum_{\nu\in\Lambda}|J_\nu(y)|^2. \tag{4}$$

It is proven in [2] that discrete least squares in multivariate polynomial spaces are stable and accurate provided that a precise condition involving m and $K(\mathbb{P}_\Lambda)$ is satisfied. Similar results have been proven in [22] for the case of noisy observations of the target function, with several noise models.

For any $\delta\in]0,1[$ we introduce the positive quantity

$$\zeta(\delta) := \delta + (1-\delta)\ln(1-\delta). \tag{5}$$

Given a threshold $\tau\in\mathbb{R}_0^+$, we introduce the truncation operator

$$T_\tau(t) := \mathrm{sign}(t)\min\{\tau,|t|\}, \qquad \text{for any } t\in\mathbb{R},$$

and use it to define the truncated discrete least-squares projection operator $u\mapsto T_\tau(\Pi_\Lambda^m u)$. The main results from [2] concerning stability and accuracy of the discrete least-squares approximation with noiseless evaluations can be summarized as follows.

Theorem 1. *In any dimension d, for any $r>0$, any $\delta\in]0,1[$ and any downward closed multi-index set $\Lambda\subset\mathbb{N}_0^d$, one has*

$$\Pr\left(\left\{(1-\delta)\|v\|^2\le\|v\|_m^2\le(1+\delta)\|v\|^2,\ \forall v\in\mathbb{P}_\Lambda\right\}\right)\ge 1-2n\exp(-\zeta(\delta)m/K(\mathbb{P}_\Lambda)). \tag{6}$$

If the following condition between m and $K(\mathbb{P}_\Lambda)$ is satisfied

$$\frac{m}{\ln m}\ge\frac{(1+r)}{\zeta(\delta)}K(\mathbb{P}_\Lambda), \tag{7}$$

then

$$\Pr\left(\left\{(1-\delta)\|v\|^2 \leq \|v\|_m^2 \leq (1+\delta)\|v\|^2,\ \forall v \in \mathbb{P}_\Lambda\right\}\right) \geq 1 - 2m^{-r}.$$

Moreover, for any $u \in L^\infty(\Gamma)$ *with* $\|u\|_{L^\infty(\Gamma)} \leq \tau$*, the following holds:*

$$\Pr\left(\|u - \Pi_\Lambda^m u\| \leq \left(1 + \sqrt{\frac{1}{1-\delta}}\right) \inf_{v \in \mathbb{P}_\Lambda} \|u - v\|_{L^\infty(\Gamma)}\right) \geq 1 - 2m^{-r}, \tag{8}$$

$$\mathbb{E}\left(\|u - T_\tau(\Pi_\Lambda^m u)\|^2\right) \leq \left(1 + \frac{4\zeta(\delta)}{(1+r)\ln m}\right)\|u - \Pi_\Lambda u\|^2 + 8\tau^2 m^{-r}. \tag{9}$$

The above theorem states that under condition (7) the discrete least-squares approximation is stable, since one has

$$(1-\delta)\|v\|^2 \leq \|v\|_m^2 \leq (1+\delta)\|v\|^2,\ \forall v \in \mathbb{P}_\Lambda \Leftrightarrow (1-\delta)\mathbf{I} \leq \mathbf{G} \leq (1+\delta)\mathbf{I},$$

where $\mathbf{I}$ denotes the identity matrix. Under the same condition, the discrete least-squares approximation is also accurate in probability, from (8), and in expectation, from (9), since the approximation error behaves like the best approximation error in L^∞ or in L^2.

The quantity $K(\mathbb{P}_\Lambda)$ depends both on Λ and on the chosen Jacobi measure, and therefore on the parameters θ_1, θ_2. The following result from [2] and [18] shows that, once these two parameters are fixed, the quantity $K(\mathbb{P}_\Lambda)$ satisfies bounds that only depend on $\#(\Lambda)$, independently of the particular shape of Λ and of the dimension d.

Lemma 1. *In any dimension* d *and for any finite downward closed set* $\Lambda \subset \mathcal{F}$*, one has*

$$\#(\Lambda) \leq K(\mathbb{P}_\Lambda) \leq \begin{cases} (\#(\Lambda))^{\ln 3/\ln 2}, & \text{if } \theta_1 = \theta_2 = -1/2, \\ (\#(\Lambda))^{2\max\{\theta_1,\theta_2\}+2} & \text{if } \theta_1, \theta_2 \in \mathbb{N}_0. \end{cases}$$

Combining the two results, one therefore obtains sufficient conditions for stability and optimal accuracy expressed only in terms of a relation between $\#(\Lambda)$ and m. For example, in the case of the uniform measure that corresponds to $\theta_1 = \theta_2 = 0$, this relation is of the form

$$\frac{m}{\ln m} \geq c\,(\#(\Lambda))^2, \quad c := c(\delta, r).$$

3 Optimal selection of downward closed polynomial spaces

The results recalled in the previous section hold for a given downward closed set $\Lambda \subset \mathcal{F}$. We now consider the problem of optimizing the choice of Λ, or equivalently that of the space $\mathbb{P}_\Lambda$.

3.1 Optimized index sets

We define the family

$$\mathcal{M}_n^d := \{\Lambda \subset \mathcal{F} \ :\ \Lambda \text{ is downward closed and } \#(\Lambda) = n\},$$

of all downward closed sets of cardinality n in d coordinates. Note that, in contrast to the family of *all* subsets of $\mathcal{F}$ of cardinality n, the family $\mathcal{M}_n^d$ is finite.

The error of best n-term polynomial approximation by downward closed sets is then defined by

$$\sigma_n(u) := \min_{\Lambda \in \mathcal{M}_n^d} \min_{v \in \mathbb{P}_\Lambda} \|u - v\|.$$

A best n-term downward closed set is a $\Lambda \in \mathcal{M}_n^d$ that achieves this minimum, that is, such that

$$\Lambda^{opt} := \operatorname*{argmin}_{\Lambda \in \mathcal{M}_n^d} \min_{v \in \mathbb{P}_\Lambda} \|u - v\|.$$

Using Parseval identity, we find that Λ^{opt} is also given by

$$\Lambda^{opt} = \operatorname*{argmin}_{\Lambda \in \mathcal{M}_n^d} \sum_{\nu \notin \Lambda} |u_\nu|^2, \quad u_\nu = \int_\Gamma u(y) J_\nu(y) d\rho(y).$$

Note that such a set may not be unique due to possible ties in the values of the coefficients, in which case we consider a unique choice by breaking the ties in some arbitrary but fixed way. We set

$$u_n := \Pi_{\Lambda^{opt}} u = \operatorname*{argmin}_{v \in \mathbb{P}_{\Lambda^{opt}}} \|u - v\|. \tag{10}$$

Of course, in the least-squares method, the discrete data do not allow us to identify Λ^{opt}. Instead, we rely on

$$\Lambda_m^{opt} := \operatorname*{argmin}_{\Lambda \in \mathcal{M}_n^d} \min_{v \in \mathbb{P}_\Lambda} \|u - v\|_m, \tag{11}$$

and compute

$$w_n := \Pi^m_{\Lambda_m^{opt}} u = \operatorname*{argmin}_{v \in \mathbb{P}_{\Lambda_m^{opt}}} \|u - v\|_m.$$

Our objective is now to compare the accuracy of the polynomial least-squares approximation based on Λ_m^{opt} with the above optimal error $\sigma_n(u)$. For this purpose, we shall use the random variable

$$C_n^d := \max_{\Lambda \in \mathcal{M}_n^d} \max_{v \in \mathbb{P}_\Lambda} \frac{\|v\|^2}{\|v\|_m^2}.$$

Note that the search of Λ_m^{opt} remains a difficult task from the computational point of view, due to the fact that $\#(\mathcal{M}_n^d)$ becomes very large even for moderate values of n and d. As we discuss further, this cardinality also affects the conditions between m and n which guarantee the optimality of the least-squares approximation based on Λ_m^{opt}.

For this reason, it is useful to introduce an additional restriction on the potential index sets. We say that Λ is *anchored* if and only if it is downward closed and satisfies in addition

$$e_j \in \Lambda \text{ and } j' \leq j \implies e_{j'} \in \Lambda,$$

where e_j and $e_{j'}$ are the Kronecker sequences with 1 at positions j and j', respectively. We also say that a polynomial space $\mathbb{P}_\Lambda$ is anchored when Λ is anchored. Likewise, we define the family

$$\mathcal{A}_n := \{\Lambda \subset \mathcal{F} \,:\, \Lambda \text{ is anchored and } \#(\Lambda) = n\}.$$

The restriction to anchored sets introduces an order of priority between the coordinates: given any $j \geq 1$, the coordinate y_j is active in Λ only if all the coordinates y_k for $k < j$ are also active. In particular, for any set $\Lambda \in \mathcal{A}_n$ we have

$$\operatorname{supp}(\Lambda) = \{1, \ldots, k\}, \tag{12}$$

for some $k \leq n - 1$.

It is proven in [8] that, for relevant classes of parametric PDEs, the same algebraic convergence rates $\mathcal{O}(n^{-s})$ can be obtained when imposing the anchored structure on the optimally selected sets $(\Lambda_n)_{n \geq 1}$ with $\#(\Lambda_n) = n$. As we shall see further, one specific advantage of anchored sets is to completely remove

the dependence on the dimension d in the convergence analysis of the least-squares method, and allows us in particular to consider the infinite-dimensional framework.

Using the same notation as before with obvious modifications, we introduce the following entities:

$$\tilde{\Lambda}^{opt} := \operatorname*{argmin}_{\Lambda\in\mathcal{A}_n} \min_{v\in\mathbb{P}_\Lambda} \|u-v\|, \qquad \tilde{u}_n := \Pi_{\tilde{\Lambda}^{opt}} u = \operatorname*{argmin}_{v\in\mathbb{P}_{\tilde{\Lambda}^{opt}}} \|u-v\|, \tag{13}$$

$$\tilde{\Lambda}_m^{opt} := \operatorname*{argmin}_{\Lambda\in\mathcal{A}_n} \min_{v\in\mathbb{P}_\Lambda} \|u-v\|_m, \qquad \tilde{w}_n := \Pi_{\tilde{\Lambda}_m^{opt}}^m u = \operatorname*{argmin}_{v\in\mathbb{P}_{\tilde{\Lambda}_m^{opt}}} \|u-v\|_m,$$

and

$$\widetilde{C}_n := \max_{\Lambda\in\mathcal{A}_n} \max_{v\in\mathbb{P}_\Lambda} \frac{\|v\|^2}{\|v\|_m^2}.$$

Remark 1. *The estimators w_n and $\tilde{w}_n$ can be viewed as the solutions of a nonconvex optimization problem. This problem has a natural algebraic formulation. Recalling $(J_\nu)_{\nu\in\mathcal{F}}$ the $L^2(\Gamma,d\rho)$-orthonormal basis, we introduce for a given finite set $\mathcal{J}\subset\mathcal{F}$ the design matrix*

$$\mathbf{D} = (J_\nu(y^i)), \tag{14}$$

where the row index i ranges in $\{1,\dots,m\}$ and the column index ν ranges in $\mathcal{J}$. Then, recalling the data vector $\mathbf{b} = (u(y^i))_{i=1,\dots,m}$, and taking $\mathcal{J}$ as the union of all downward closed sets of cardinality n, we find that the component vector $\mathbf{a} = (a_\nu)_{\nu\in\mathcal{J}}$ of $w_n = \sum_{\nu\in\mathcal{J}} a_\nu J_\nu$ is the solution to the constrained minimization

$$\min\{\|\mathbf{D}\mathbf{a}-\mathbf{b}\|_{\ell^2} \ : \ \|\mathbf{a}\|_{\ell_d^0} \le n\}.$$

Here $\|\mathbf{a}\|_{\ell_d^0}$ is the cardinality of the smallest downward closed set $\Lambda\subset\mathcal{J}$ that contains all the nonzero entries of $\mathbf{a}$. In other words, we minimize over those $\mathbf{a}$ whose support is contained in a downward closed set of cardinality at most n. Likewise, taking $\mathcal{J}$ as the union of all anchored sets of cardinality n, we find that the component vector $\widetilde{\mathbf{a}} = (\tilde{a}_\nu)_{\nu\in\mathcal{J}}$ of $\tilde{w}_n = \sum_{\nu\in\mathcal{J}} \tilde{a}_\nu J_\nu$ is the solution of the constrained minimization

$$\min\{\|\mathbf{D}\widetilde{\mathbf{a}}-\mathbf{b}\|_{\ell^2} \ : \ \|\widetilde{\mathbf{a}}\|_{\ell_a^0} \le n\}.$$

Here $\|\widetilde{\mathbf{a}}\|_{\ell_a^0}$ is the cardinality of the smallest anchored set $\Lambda\subset\mathcal{J}$ that contains all the nonzero entries of $\widetilde{\mathbf{a}}$. It is well known that such optimization problems with ℓ^0-type constraints have combinatorial nature, and in turn numerical algorithms for computing their solutions do not generally have polynomial complexity. This justifies in practice the use of convex relaxation methods such as basis pursuit, or greedy selection strategies such as orthonormal matching pursuit. While our paper is mainly directed towards the convergence properties of the estimators w_n and $\tilde{w}_n$ which are the exact solutions to the above problems, we discuss these numerical methods in the final section.

We now would like to compare the estimators w_n and $\tilde{w}_n$ with the best n-term approximation in the aforementioned classes of multi-index sets $\mathcal{M}_n^d$ and $\mathcal{A}_n$. The following lemma shows the role played by C_n^d and $\widetilde{C}_n$ in quantifying the relation between the error achieved by the optimal discrete least-squares projection and the error achieved by the optimal L^2 projection.

Lemma 2. *It holds that, for any $\Lambda\in\mathcal{M}_n^d$,*

$$\|u-w_n\| \le \|u-v\| + 2\sqrt{C_{2n-1}^d}\|u-v\|_m, \quad v\in\mathbb{P}_\Lambda, \tag{15}$$

and for any $\Lambda\in\mathcal{A}_n$,

$$\|u-\tilde{w}_n\| \le \|u-\tilde{v}\| + 2\sqrt{\widetilde{C}_{2n-1}}\|u-\tilde{v}\|_m, \quad \tilde{v}\in\mathbb{P}_\Lambda. \tag{16}$$

Proof. Let $\Lambda \in \mathcal{M}_n^d$ and define $\widehat{\Lambda} := \Lambda \cup \Lambda_m^{opt}$. We first observe that $\widehat{\Lambda}$ is also downward closed and $\#(\widehat{\Lambda}) \leq 2n-1$ because any downward closed set contains the null multi-index. Since $w_n \in \mathbb{P}_{\Lambda_m^{opt}}$, we have $v - w_n \in \mathbb{P}_{\widehat{\Lambda}}$ for any $v \in \mathbb{P}_\Lambda$. It follows that

$$\|v - w_n\| \leq \sqrt{C_{2n-1}^d}\|v - w_n\|_m \leq \sqrt{C_{2n-1}^d}(\|u - v\|_m + \|u - w_n\|_m) \leq 2\sqrt{C_{2n-1}^d}\|u - v\|_m,$$

and therefore

$$\|u - w_n\| \leq \|u - v\| + \|v - w_n\| \leq \|u - v\| + 2\sqrt{C_{2n-1}^d}\|u - v\|_m,$$

which is (15). The proof of (16) is analogous. □

Note that the estimates in the above lemma imply in particular that

$$\|u - w_n\| \leq \|u - u_n\| + 2\sqrt{C_{2n-1}^d}\|u - u_n\|_m, \tag{17}$$

and

$$\|u - \tilde{w}_n\| \leq \|u - \tilde{u}_n\| + 2\sqrt{\widetilde{C}_{2n-1}}\|u - \tilde{u}_n\|_m,$$

with u_n and $\tilde{u}_n$ defined by (10) and (13). Note that they also imply

$$\|u - w_n\| \leq \left(1 + 2\sqrt{C_{2n-1}^d}\right)\|u - v\|_{L^\infty}, \quad v \in \mathbb{P}_\Lambda, \tag{18}$$

for any $\Lambda \in \mathcal{M}_n^d$, and

$$\|u - \tilde{w}_n\| \leq \left(1 + 2\sqrt{\widetilde{C}_{2n-1}}\right)\|u - \tilde{v}\|_{L^\infty}, \quad \tilde{v} \in \mathbb{P}_\Lambda, \tag{19}$$

for any $\Lambda \in \mathcal{A}_n$.

3.2 Probabilistic bounds

In view of Lemma 2, we are interested in bounding the random variables C_n^d and $\widetilde{C}_n$. In this section we give probabilistic bounds, which ensure that under certain conditions between m and n, these random variables do not exceed a fixed value, here set to 2, with high probability. In the whole section we choose $\delta = 1/2$, so that, with the notation (5), one has

$$\zeta := \zeta(\delta) = \zeta(1/2) = (1 - \ln 2)/2 \approx 0.153.$$

We define, for any $\nu \in \mathcal{F}$, the "rectangular" set $\mathcal{R}_\nu := \{\mu \in \mathcal{F},\ \mu \leq \nu\}$, and for any $n \geq 1$, the hyperbolic cross set

$$\mathcal{H}_n^d := \left\{\mu \in \mathcal{F} :\ \prod_{j=1}^d (\mu_j + 1) \leq n\right\}.$$

Note that

$$\mathcal{H}_n^d = \bigcup_{\#(\mathcal{R}_\nu) \leq n} \mathcal{R}_\nu.$$

The cardinality of $\mathcal{H}_n^d$ is bounded by

$$\#(\mathcal{H}_n^d) \leq n(1 + \ln(n))^{d-1}, \tag{20}$$

see [17, Appendix A.2] for a proof and some remarks on the accuracy of this upper bound. The relevance of the hyperbolic cross for our purposes is due to the following observation.

Lemma 3. *The union of all downward closed sets of cardinality at most n in finite dimension d coincides with $\mathcal{H}_n^d$, that is,*

$$\bigcup_{\Lambda \in \mathcal{M}_n^d} \Lambda = \mathcal{H}_n^d. \tag{21}$$

Proof. On the one hand, all rectangles $\mathcal{R}_\nu$ such that $\#(\mathcal{R}_\nu) \leq n$ belong to $\mathcal{M}_n^d$, so that inclusion holds from right to left. On the other hand, inclusion from left to right follows by observing that for any $\Lambda \in \mathcal{M}_n^d$, one has $\Lambda = \cup_{\mu \in \Lambda} \mathcal{R}_\mu$ and $\mathcal{R}_\mu \subset \mathcal{H}_n^d$ for all $\mu \in \Lambda$. □

This leads us to a first probabilistic bound for the random variable C_n^d. Indeed, using (21) we obtain that

$$\Pr\left(C_n^d > 2\right) = \Pr\left(\max_{\Lambda \in \mathcal{M}_n^d} \max_{v \in \mathbb{P}_\Lambda} \frac{\|v\|^2}{\|v\|_m^2} > 2\right) \leq \Pr\left(\max_{v \in \mathbb{P}_{\mathcal{H}_n^d}} \frac{\|v\|^2}{\|v\|_m^2} > 2\right).$$

Thus, using Theorem 1 with $\delta = 1/2$ combined with the estimates in Lemma 1, we obtain that, in any dimension d and for any $r > 0$, if m and n satisfy

$$\frac{m}{\ln m} \geq \begin{cases} \frac{(1+r)}{\zeta}(\#(\mathcal{H}_n^d))^{\ln 3/\ln 2}, & \text{with Chebyshev polynomials of the first kind,} \\ \frac{(1+r)}{\zeta}(\#(\mathcal{H}_n^d))^{2\max\{\theta_1,\theta_2\}+2}, & \text{with Jacobi polynomials and } \theta_1, \theta_2 \in \mathbb{N}_0, \end{cases} \tag{22}$$

then

$$\Pr(C_n^d > 2) \leq 2m^{-r}.$$

From (21) and (12) we also find that the union of all anchored sets of cardinality at most n satisfies the following inclusion

$$\bigcup_{\Lambda \in \mathcal{A}_n} \Lambda \subset \mathcal{H}_n^{n-1}.$$

By similar arguments, we obtain the following probabilistic bound for the random variable $\widetilde{C}_n$: in any dimension d, for any $r > 0$, if m and n satisfy

$$\frac{m}{\ln m} \geq \begin{cases} \frac{(1+r)}{\zeta}(\#(\mathcal{H}_n^{n-1}))^{\ln 3/\ln 2}, & \text{with Chebyshev polynomials of the first kind,} \\ \frac{(1+r)}{\zeta}(\#(\mathcal{H}_n^{n-1}))^{2\max\{\theta_1,\theta_2\}+2}, & \text{with Jacobi polynomials and } \theta_1, \theta_2 \in \mathbb{N}_0, \end{cases} \tag{23}$$

then

$$\Pr(\widetilde{C}_n > 2) \leq 2m^{-r}.$$

The above results describe the regimes of m and n such that C_n^d and $\widetilde{C}_n$ do not exceed 2 with high probability. In the case of downward closed sets, this regime is quite restrictive due to the presence of $\ln(n)^{d-1}$ in the upper bound (20) for the cardinality of $\mathcal{H}_n^d$, which enforces the sample size m to be extremely large as d grows. Likewise, m has to be extremely large compared to n in the case of anchored sets.

We next describe another strategy which yields similar probabilistic bounds under less restrictive regimes. It is based on estimating the cardinality of $\mathcal{M}_n^d$ and $\mathcal{A}_n$ and using union bounds. Our way of estimating $\#(\mathcal{M}_n^d)$ and $\#(\mathcal{A}_n)$ is based on strategies for encoding any downward closed or anchored set. One first such strategy leads to the following result.

Lemma 4 (First cardinality bound)**.** *One has the cardinality estimates*

$$\#(\mathcal{M}_n^d) \leq 2^{nd}, \tag{24}$$

and

$$\#(\mathcal{A}_n) \leq 2^{n(n-1)}. \tag{25}$$

Proof. We encode any downward closed set $\Lambda \subset \mathcal{F}$ of cardinality n in d dimensions, by associating d bits $b_{\nu,1}, \dots, b_{\nu,d}$ to each multi-index $\nu \in \Lambda$, where the value of the jth bit is equal to one if $\nu + e_j \in \Lambda$, and equal to zero if $\nu + e_j \notin \Lambda$. We order these blocks of bits according to the lexicographic order $\nu^1, \dots, \nu^n$ of appearance of ν in Λ, where $\nu^1 = (0, \dots, 0)$. The resulting bitstream

$$B_\Lambda := (b_{\nu^1,1}, \dots, b_{\nu^1,d}, b_{\nu^2,1}, \dots, b_{\nu^2,d}, \dots, b_{\nu^n,1}, \dots, b_{\nu^n,d}),$$

uniquely encodes Λ, that is, the encoding map $\Lambda \mapsto B_\Lambda$ is injective. Indeed, assuming that $\nu^1, \dots, \nu^k$ have been already identified, then the partial bitstream $(b_{\nu^1,1}, \dots, b_{\nu^k,d})$ contains the information on the positions of all indices $\nu \in \Lambda$ such that $\nu \notin \{\nu^1, \dots, \nu^k\}$ and $\nu - e_j \in \{\nu^1, \dots, \nu^k\}$ for some j. Therefore ν^{k+1} is uniquely determined as the index with smallest lexicographic order among such indices.

Since the length of B_Λ is nd, this model leads us to the upper bound (24). The same encoding can be applied to anchored sets of cardinality at most n with $n-1$ bits for each index, leading to (25), which also directly follows from (24) and the fact that $\#(\mathcal{A}_n) \leq \#(\mathcal{M}_n^{n-1})$. □

Remark 2. *Note that the cardinality estimate for the anchored sets is independent of the dimension d. In particular, this allows us to derive some further results in the infinite-dimensional framework, when using anchored sets.*

Recalling the definition of $K(\mathbb{P}_\Lambda)$ from (4), we introduce the following notation:

$$K_n = \max_{\Lambda \in \mathcal{M}_n^d} K(\mathbb{P}_\Lambda),$$

$$\widetilde{K}_n = \max_{\Lambda \in \mathcal{A}_n} K(\mathbb{P}_\Lambda).$$

Note that, according to Lemma 1, one has the estimate

$$\widetilde{K}_n \leq K_n \leq \begin{cases} n^{\ln 3/\ln 2}, & \text{if } \theta_1 = \theta_2 = -1/2, \\ n^{2\max\{\theta_1,\theta_2\}+2} & \text{if } \theta_1, \theta_2 \in \mathbb{N}_0. \end{cases} \tag{26}$$

We are now in position to establish probabilistic bounds for the random variable C_n^d and $\widetilde{C}_n$.

Lemma 5. *In any finite dimension d, for any $r > 0$, if m and n satisfy*

$$\frac{m}{\ln m} \geq \left(1 + r + \frac{nd \ln 2}{\ln m}\right) \frac{K_n}{\zeta}, \tag{27}$$

then

$$\Pr\left(C_n^d > 2\right) \leq 2nm^{-(r+1)} \leq 2m^{-r}.$$

In any finite dimension d or in infinite dimension, for any $r > 0$, if m and n satisfy

$$\frac{m}{\ln m} \geq \left(1 + r + \frac{n^2 \ln 2}{\ln m}\right) \frac{\widetilde{K}_n}{\zeta}, \tag{28}$$

then

$$\Pr\left(\widetilde{C}_n > 2\right) \leq 2nm^{-(r+1)} \leq 2m^{-r}.$$

Proof. Recalling (6) and using a union bound, we obtain

$$\begin{aligned}\Pr\left(C_n^d > 2\right) &= \Pr\left(\max_{\Lambda\in\mathcal{M}_n^d}\max_{v\in\mathbb{P}_\Lambda}\frac{\|v\|^2}{\|v\|_m^2} > 2\right)\\ &\le \sum_{\Lambda\in\mathcal{M}_n^d}\Pr\left(\max_{v\in\mathbb{P}_\Lambda}\frac{\|v\|^2}{\|v\|_m^2} > 2\right)\\ &\le 2^{nd}2n\exp\left\{-\zeta m/K_n\right\}\\ &= 2n\exp\left\{-\zeta m/K_n + nd\ln(2)\right\},\end{aligned}$$

where we have used the cardinality estimate (24). The final bound is smaller than $2n^{-r}$ under condition (27). The proof for $\widetilde{C}_n$ is completely similar, using the cardinality estimate (25). □

Remark 3. *Combining with* (26)*, we find that* (27) *holds if* r*,* d*,* m *and* n *satisfy*

$$\frac{m}{\ln m} \ge \begin{cases}\left(1+r+\dfrac{nd\ln 2}{\ln m}\right)\dfrac{n^{\ln 3/\ln 2}}{\zeta}, & \text{with Chebyshev polynomials of the first kind,}\\ \left(1+r+\dfrac{nd\ln 2}{\ln m}\right)\dfrac{n^{2\max\{\theta_1,\theta_2\}+2}}{\zeta}, & \text{with Jacobi polynomials and } \theta_1,\theta_2\in\mathbb{N}_0.\end{cases} \tag{29}$$

Similarly, we find that (28) *holds if* r*,* m *and* n *satisfy*

$$\frac{m}{\ln m} \ge \begin{cases}\left(1+r+\dfrac{n^2\ln 2}{\ln m}\right)\dfrac{n^{\ln 3/\ln 2}}{\zeta}, & \text{with Chebyshev polynomials of the first kind,}\\ \left(1+r+\dfrac{n^2\ln 2}{\ln m}\right)\dfrac{n^{2\max\{\theta_1,\theta_2\}+2}}{\zeta}, & \text{with Jacobi polynomials and } \theta_1,\theta_2\in\mathbb{N}_0.\end{cases} \tag{30}$$

We next give an improved result on the cardinality of $\mathcal{M}_n^d$ and $\mathcal{A}_n$ based on another encoding strategy.

Lemma 6 (Second cardinality bound)**.** *One has the cardinality estimates*

$$\#(\mathcal{M}_n^d) \le d^{n-1}(n-1)!, \tag{31}$$

and

$$\#(\mathcal{A}_n) \le \left((n-1)!\right)^2. \tag{32}$$

Proof. Given any downward closed multi-index set Λ with $\#(\Lambda) = n$, we order the elements of Λ in such a way that the set

$$\Lambda^k := \{\nu^1,\dots,\nu^k\},$$

obtained by retaining only the first k elements of Λ, is downward closed for any $k = 1,\dots,n$. Such an ordering always exists, and in general it is not unique. Notice that it always holds $\nu^1 = (0,\dots,0)$. One way to impose a unique ordering is by taking for ν^k the smallest index ν in lexicographic order among those $\nu \in \Lambda\setminus\{\nu^1,\dots,\nu^{k-1}\}$ such that $\{\nu^1,\dots,\nu^k\}$ is downward closed. Each ν^k can be uniquely characterized by choosing some $l_k \in \{1,\dots,k-1\}$ and $j_k \in \{1,\dots,d\}$ such that

$$\nu^k = \nu^{l_k} + e_{j_k}.$$

Again this choice can be made unique by asking that j_k is the smallest number with such a property. Therefore, the resulting map

$$\Lambda \mapsto (j_2, l_3, j_3, \dots, l_n, j_n),$$

is well defined and injective. Hence

$$\#(\mathcal{M}_n^d) \le d(2d)\cdots(n-1)d,$$

which is (31). We obtain (32) from the inequality $\#(\mathcal{A}_n) \le \#(\mathcal{M}_n^{n-1})$. □

The above results lead to improved probabilistic bounds for the random variables C_n^d and $\widetilde{C}_n$.

Lemma 7. *In any finite dimension d, for any $r > 0$, if m and n satisfy*

$$\frac{m}{\ln m} \geq \left(1 + r + \frac{n \ln(dn)}{\ln m}\right) \frac{K_n}{\zeta}, \tag{33}$$

then

$$\Pr\left(C_n^d > 2\right) \leq 2nm^{-(r+1)} \leq 2m^{-r}.$$

In any finite dimension d or in infinite dimension, for any $r > 0$, if m and n satisfy

$$\frac{m}{\ln m} \geq \left(1 + r + \frac{2n \ln n}{\ln m}\right) \frac{\widetilde{K}_n}{\zeta}, \tag{34}$$

then

$$\Pr\left(\widetilde{C}_n > 2\right) \leq 2nm^{-(r+1)} \leq 2m^{-r}.$$

Proof. Recalling (6) and using the inequality $n! \leq e\sqrt{n}(n/e)^n$ which holds for any $n \geq 1$, we obtain by means of a union bound that, for any $n \geq 2$,

$$\begin{aligned}
\Pr\left(C_n^d > 2\right) &\leq d^{n-1}(n-1)!\ (2n)\ \exp\left\{-\zeta m/K_n\right\} \\
&\leq d^{n-1} e\sqrt{n-1}\left(\frac{n-1}{e}\right)^{n-1}\ (2n)\ \exp\left\{-\zeta m/K_n\right\} \\
&= 2n \exp\left\{-\zeta m/K_n + (n-1/2)\ln\left(\frac{d(n-1)}{e}\right) - \frac{1}{2}\ln\left(\frac{d}{e}\right) + 1\right\} \\
&\leq 2n \exp\left\{-\zeta m/K_n + n\ln(dn)\right\},
\end{aligned}$$

where we have used the cardinality estimate (31). The final bound is smaller than $2n^{-r}$ under condition (33). Trivially the final bound holds true also when $n = 1$. The proof for $\widetilde{C}_n$ is completely similar, using the cardinality estimate (32). □

Remark 4. *Combining with* (26)*, we find that* (33) *holds if r, d, m and n satisfy*

$$\frac{m}{\ln m} \geq \begin{cases} \left(1 + r + \dfrac{n \ln(dn)}{\ln m}\right) \dfrac{n^{\ln 3/\ln 2}}{\zeta}, & \text{with Chebyshev polynomials of the first kind,} \\ \left(1 + r + \dfrac{n \ln(dn)}{\ln m}\right) \dfrac{n^{2\max\{\theta_1,\theta_2\}+2}}{\zeta}, & \text{with Jacobi polynomials and } \theta_1, \theta_2 \in \mathbb{N}_0. \end{cases} \tag{35}$$

Similarly, we find that (34) *holds if r, m and n satisfy*

$$\frac{m}{\ln m} \geq \begin{cases} \left(1 + r + \dfrac{2n \ln n}{\ln m}\right) \dfrac{n^{\ln 3/\ln 2}}{\zeta}, & \text{with Chebyshev polynomials of the first kind,} \\ \left(1 + r + \dfrac{2n \ln n}{\ln m}\right) \dfrac{n^{2\max\{\theta_1,\theta_2\}+2}}{\zeta}, & \text{with Jacobi polynomials and } \theta_1, \theta_2 \in \mathbb{N}_0. \end{cases} \tag{36}$$

The regimes of m and n described by the above results are in principle less restrictive than those previously obtained using the cardinality of $\mathcal{H}_n^d$ or $\mathcal{H}_n^{n-1}$. Indeed, for most regimes of n and d the cardinalities $\#(\mathcal{H}_n^d)$ or $\#(\mathcal{H}_n^{n-1})$ are larger than $n\ln(dn)$ and $n\ln(n)$, respectively. We may summarize the probabilistic bounds established in this section as follows: for any $r > 0$ and any $n \geq 1$ one has $C_n^d \leq 2$ with probability larger than $1 - 2m^{-r}$ provided that (22) or (29) or (35) holds. Likewise, one has $\widetilde{C}_n \leq 2$ with probability larger than $1 - 2m^{-r}$ provided that (23) or (30) or (36) holds.

Remark 5. *The encoding strategies that are used for proving Lemma 4 and Lemma 6 are both redundant, leading to an overestimation of $\#(\mathcal{M}_n^d)$ and $\#(\mathcal{A}_n)$. We are not aware of estimates for these cardinalities which are provably sharp up to multiplicative constants independent of n and d. However, we can establish lower bounds which show that for certain particular regimes of n and d, these cardinalities grow exponentially fast. One simple instance of a lower bound for $\#(\mathcal{M}_n^d)$ in the regime where $n-1 \le d$ is obtained as follows: we note that $\mathcal{M}_n^d$ includes in particular all sets of the form $\{(0,\dots,0)\} \cup \{e_j \,:\, j \in S\}$ for $S \subset \{1,\dots,d\}$ such that $\#(S) = n-1$. It follows that*

$$\#(\mathcal{M}_n^d) \ge \binom{d}{n-1}.$$

In the regime where $n = d/2$ (for even d), we thus find that $\log_2(\#(\mathcal{M}_n^d))$ grows at least as fast as d.

Remark 6. *It is interesting to compare the probabilistic bounds obtained in this section with Restricted Isometry Properties (RIP) recently obtained in [5], that are a common vehicle in the analysis of compressed sensing schemes. Recalling the design matrix $\mathbf{D}$ introduced in (14), and defining its renormalized version $\Phi := m^{-1/2}\mathbf{D}$, we indeed see that the property $C_n^d \le 2$ can be rephrased as*

$$\frac{1}{2}\|\mathbf{a}\|^2 \le \|\Phi\mathbf{a}\|^2, \quad \|\mathbf{a}\|_{\ell_d^0} \le n,$$

that is, for all $\mathbf{a}$ whose support is contained in a downward closed set of cardinality at most n. In [5], it is shown that the RIP property

$$(1-\delta)\|\mathbf{a}\|^2 \le \|\Phi\mathbf{a}\|^2 \le (1+\delta)\|\mathbf{a}\|^2, \quad \|\mathbf{a}\|_{\ell_d^0} \le n, \tag{37}$$

holds with probability at least $1-\gamma$ if

$$m \ge 2^6 e \frac{K_n}{\tilde\delta^2}\ln\Big(\frac{K_n}{\tilde\delta^2}\Big)\max\Big\{\frac{2^5}{\tilde\delta^4}\ln\Big(40\frac{K_n}{\tilde\delta^2}\ln\Big(\frac{K_n}{\tilde\delta^2}\Big)\Big)\ln(4N), \frac{1}{\tilde\delta}\ln\Big(\frac{1}{\gamma\tilde\delta}\ln\Big(\frac{K_n}{\tilde\delta^2}\Big)\Big)\Big\}, \quad \tilde\delta := \frac{\delta}{13}, \tag{38}$$

where $N = \#(\mathcal{J})$ with $\mathcal{J}$ the union of all downward closed sets of cardinality at most n, that is, $\mathcal{J} = \mathcal{H}_n^d$. Note that in our analysis, we are only interested in establishing the lower inequality in the RIP, with the particular constant $\delta = \frac{1}{2}$. However, since we have started from the two-sided estimate in (6), one easily checks that the same analysis leads to the validity of the RIP property (37) with probability at least $1-\gamma$, under the regime

$$\frac{m}{\ln m} \ge \left(1 + \frac{\ln(2/\gamma)}{\ln m} + \frac{n\ln(dn)}{\ln m}\right)\frac{K_n}{\zeta(\delta)}, \tag{39}$$

where $\zeta(\delta)$ is again given by (5). From an asymptotic point of view, it can be checked that the regime (38) is more favorable than (39): indeed n appears on the right side of (38) only through K_n up to logarithmic factors, while it appears through nK_n in the right side of (39). However, the multiplicative constant on the right side of (38) is prohibitively large: for example in the case $\delta = \frac{1}{2}$ which is considered in the present paper, we find that $\tilde\delta = \frac{1}{26}$ and therefore $2^6 e \frac{1}{\tilde\delta^2}\frac{2^5}{\tilde\delta^4} > 10^{12}$. This shows that the regime described by (39) is more favorable for moderate values of n. Similar remarks hold when considering anchored sets rather than downward closed sets.

3.3 Accuracy of the optimized discrete least-squares approximation

We are now in position to state our main results concerning the accuracy of the discrete least-squares approximation w_n and $\tilde w_n$ over the optimized index set Λ_m^{opt} and $\tilde\Lambda_m^{opt}$. These results show that the accuracy compares favorably with the best approximation error of the function u, measured either in L^∞ or L^2, using optimal choices of downward closed or anchored sets (which might differ from the sets Λ_m^{opt} and $\tilde\Lambda_m^{opt}$). We begin with a result expressed in probability.

Theorem 2. *Consider a function u defined on Γ and let $r > 0$. In any finite dimension, under condition* (22)*, or* (29)*, or* (35)*, with n replaced by $2n-1$, it holds that*

$$\Pr\left(\|u - w_n\| \leq (1+2\sqrt{2}) \min_{\Lambda \in \mathcal{M}_n^d} \min_{v \in \mathbb{P}_\Lambda} \|u - v\|_{L^\infty(\Gamma)}\right) \geq 1 - 2m^{-r}.$$

In any finite or infinite dimension, under condition (23)*, or* (30)*, or* (36)*, with n replaced by $2n-1$, it holds that*

$$\Pr\left(\|u - \tilde{w}_n\| \leq (1+2\sqrt{2}) \min_{\Lambda \in \mathcal{A}_n} \min_{v \in \mathbb{P}_\Lambda} \|u - v\|_{L^\infty(\Gamma)}\right) \geq 1 - 2m^{-r}.$$

Proof. These estimates immediately follow from (18) and (19) combined with the probabilistic bounds from the previous section. □

We next give a result expressed in expectation for the truncated discrete least-squares projection $T_\tau(w_n)$ and $T_\tau(\tilde{w}_n)$.

Theorem 3. *Consider a function u defined on Γ, such that $|u(y)| \leq \tau$ for any $y \in \Gamma$, and let $r > 0$. In any finite dimension, under condition* (22)*, or* (29)*, or* (35)*, with n replaced by $2n-1$, it holds that*

$$\mathbb{E}(\|u - T_\tau(w_n)\|^2) \leq (9+4\sqrt{2})\|u - u_n\|^2 + 8\tau^2 m^{-r}. \tag{40}$$

In any finite or infinite dimension, under condition (23)*, or* (30)*, or* (36)*, with n replaced by $2n-1$, it holds that*

$$\mathbb{E}(\|u - T_\tau(\tilde{w}_n)\|^2) \leq (9+4\sqrt{2})\|u - \tilde{u}_n\|^2 + 8\tau^2 m^{-r}. \tag{41}$$

Proof. For (40), we distinguish between the two complementary events $\Omega_1 := \{C_{2n-1}^d \leq 2\}$ and $\Omega_2 := \{C_{2n-1}^d > 2\}$ and write

$$\mathbb{E}(\|u - T_\tau(w_n)\|^2) = \mathbb{E}(\|u - T_\tau(w_n)\|^2 \,|\, \Omega_1) \Pr(\Omega_1) + \mathbb{E}(\|u - T_\tau(w_n)\|^2 \,|\, \Omega_2) \Pr(\Omega_2) =: E_1 + E_2.$$

Since $|u - T_\tau(w_n)| \leq 2\tau$ and $\Pr(\Omega_2) \leq 2m^{-r}$, the second term E_2 is bounded by $8\tau^2 m^{-r}$. For the first term E_1, we combine (17) and the fact that $|u - T_\tau(w_n)| \leq |u - w_n|$ to obtain the bound

$$\begin{aligned} E_1 \leq & \mathbb{E}\Big((\|u - u_n\| + 2\sqrt{2}\|u - u_n\|_m)^2\Big) \\ = & \|u - u_n\|^2 + 4\sqrt{2}\|u - u_n\|\mathbb{E}\Big(\|u - u_n\|_m\Big) + 8\mathbb{E}\Big(\|u - u_n\|_m^2\Big) \\ \leq & (9+4\sqrt{2})\|u - u_n\|^2. \end{aligned}$$

The proof of (41) is analogous. □

Remark 7. *The constants $1+2\sqrt{2}$ and $9+4\sqrt{2}$ in the above theorems can be reduced if one further restricts the regime between m and n so that C_{2n-1}^d and $\tilde{C}_{2n-1}$ are close to 1 with high probability.*

4 Alternative algorithms for model selection

The actual computation of the optimized discrete least-squares approximations w_n and $\tilde{w}_n$ in Theorems 2 and 3 would require an exhaustive search among all possible choices of downward closed or anchored sets of a given cardinality n, and this task might become computationally prohibitive when n and d are simultaneously large. Our results should therefore mainly be viewed as a benchmark in arbitrary dimension d for assessing the performance of fast selection algorithms.

We review hereafter other strategies, which could be used for the purpose of selecting a proper polynomial space, and relate them with the results from this paper.

- Model selection by complexity regularization: the above discussed least-squares methods on optimized downward closed or anchored sets may be viewed as an instance of a model selection procedure. Model selection is a widely studied topic in statistical learning theory, in various settings such as classification, density estimation, denoising, regression and inverse problems. In the regression framework, a typical approach consists in adding a complexity penalization term to the least-squares empirical risk to be minimized, see for example Chapter 12 in [14]. One general result for such estimators is provided by Theorem 12.1 of [14] established in the bounded regression framework. It gives an oracle bound which has the typical form of the minimum over all considered models of the sum of the approximation error and of a penalty term which is always larger than m^{-1}, and which persists even as the noise level tends to zero. Therefore we cannot retrieve from these standard results for model selection the potentially fast convergence rates which can be derived from the above Theorem 3 in the noisefree case. Let us note that, from a computational point of view, this approach has the same complexity as the one required to compute the estimators w_n or $\tilde{w}_n$ in the present paper, since it is based on an exhaustive optimization over the index sets. It is therefore prohibitive for large values of d and n.

- Convex relaxation of ℓ^0 problems: as explained in Remark 1, the estimators w_n and $\tilde{w}_n$ are solutions to nonconvex optimization problems of ℓ^0 type up the additional prescription of the downward closed or anchored structure of the index sets. Convex relaxation methods based on ℓ^1 or weighted-ℓ^1 minimization, such as basis pursuit or LASSO, have been intensively studied, in particular under RIP properties for the design matrix. In the context of multivariate approximation, they have been first studied in the trigonometric polynomial framework [9], and then in the algebraic polynomial framework [24, 25, 5]. The corresponding optimization algorithms are computationally much less intensive than the complete optimization of the index set which is needed to compute w_n or $\tilde{w}_n$, yet their complexity is still polynomial in the cardinality of the set $\mathcal{J}$ that indexes the columns of $\mathbf{D}$. This set coincides with the hyperbolic cross $\mathcal{H}_n^d$ in the downard closed case, leading to a computational cost that could still be prohibitive for simultaneously large values of n and d. Note that these methods do not necessarily generate downward closed or anchored sets. While one may use the compressed sensing theory based on RIP properties to analyze the accuracy of the resulting estimators, see in particular the recovery guarantee established in Theorem 4.7 of [5], it is not clear to us that they achieve the same optimality bounds as described by Theorem 2 and 3.

- Greedy algorithms: one classical alternative to the above described convex relaxation methods are greedy algorithms such as orthonormal matching pursuit (OMP) and its variants, such as iterative hard or soft thresholding. Convergence bounds for the estimators produced by theses algorithms have been established under RIP properties, see [6, 28] for OMP in a general framework, and Theorem 4.9 of [5] for iterative hard thresholding in the context of multivariate polynomial approximation. Similar to convex relaxation methods, it is not clear that the estimators obtained by these approaches achieve the same optimality bounds as described by Theorem 2 and 3. From a computational point of view, the complexity of each step of OMP scales linearly in the cardinality of $\mathcal{J}$ leading to a smaller computational cost than convex relaxation methods, yet that could still be prohibitive for simultaneously large values of n and d. One natural way to reduce the complexity is to enforce the downward closed or anchored structure in the index selection: if Λ_k is the index set generated after k steps of OMP, we construct $\Lambda_{k+1} = \Lambda_k \cup \{\nu\}$ by maximizing the inner product of the residual with the colums of $\mathbf{D}$ corresponding only to the indices ν such that the set $\Lambda_k \cup \{\nu\}$ remains downward closed (or anchored). This restriction has the effect of significantly reducing the complexity. However it is not clear that any recovery guarantee can be established for such an algorithm.

- Relaxed minimization: let us observe that one could replace the set Λ_m^{opt} defined in (11) with a near-optimal set Λ_m^{near} in the sense that one has

$$\min_{v\in\mathbb{P}_{\Lambda_m^{near}}} \|u-v\|_m \leq C \min_{\Lambda\in\mathcal{M}_n^d} \min_{v\in\mathbb{P}_\Lambda} \|u-v\|_m,$$

for some fixed constant $C \geq 1$. Then, it is easily checked that similar convergence bounds can be established for the resulting estimators, up to changing the multiplicative constants. If one is only interested in establishing convergence rates, the optimality criterion can be even further relaxed by only asking that

$$\min_{v\in\mathbb{P}_{\Lambda_m^{near}}} \|u-v\|_m \leq C \min_{\Lambda\in\mathcal{M}_{\xi n}^d} \min_{v\in\mathbb{P}_\Lambda} \|u-v\|_m,$$

for some fixed $0 < \xi \leq 1$. However, designing a fast selection algorithm that would produce such near-optimal sets is currently an open problem. A similar objective has been accomplished in the setting of tree structured index sets, see [1].

References

[1] P. Binev and R. DeVore, *Fast computation in adaptive tree approximation*, Numer. Math., 97:193–217, 2004.

[2] A. Chkifa, A. Cohen, G. Migliorati, F. Nobile, and R. Tempone: *Discrete least squares polynomial approximation with random evaluations - application to parametric and stochastic elliptic PDEs*, ESAIM Math. Model. Numer. Anal., 49(3):815–837, 2015.

[3] A. Chkifa, A. Cohen, and C. Schwab: *High-dimensional adaptive sparse polynomial interpolation and applications to parametric PDEs*, Found. Comput. Math., 14:601–633, 2014.

[4] A. Chkifa, A. Cohen, and C. Schwab: *Breaking the curse of dimensionality in parametric and stochastic PDEs*, J. Math. Pures Appl., 103(2):400–428, 2015.

[5] A. Chkifa, N. Dexter, H. Tran, and C. Webster: *Polynomial approximation via compressed sensing of high-dimensional functions on lower sets*, arXiv:1602.05823, 2016.

[6] A. Cohen, W. Dahmen, and R. DeVore, Orthogonal matching pursuit under the restricted isometry property, to appear in Constr. Approx., 2016.

[7] A. Cohen, M.A. Davenport, and D. Leviatan: *On the stability and accuracy of least squares approximations*, Found. Comput. Math., 13:819–834, 2013.

[8] A. Cohen and R. DeVore: *Approximation of high-dimensional parametric PDEs*, Acta Numer., 24:1–159, 2015.

[9] A. Cohen, R. DeVore, S. Foucart, and H. Rauhut, *Recovery of functions of many variables via compressed sensing*, Proceedings SampTA11, Singapore, 2011.

[10] A. Cohen, R. DeVore, and C. Schwab: *Analytic regulality and polynomial approximation of parametric and stochastic PDEs*, Anal. Appl., 9:11–47, 2011.

[11] F. Cucker and D. Zhou: *Learning Theory: An Approximation Theory Viewpoint*, Cambridge University Press, 2007.

[12] C. de Boor and A. Ron: *Computational aspects of polynomial interpolation in several variables*, Math. Comp., 58:705–727, 1992.

[13] N. Dyn and M.S. Floater: *Multivariate polynomial interpolation on lower sets*, J. Approx. Theory, 177:3–42, 2014.

[14] L. Györfi, M. Kohler, A. Krzyzak, and H.Walk: *A distribution-free theory of nonparametric regression*, Springer-Verlag, 2002.

[15] J. Kuntzman: *Méthodes numériques - Interpolation, dérivées*, Dunod, Paris, 1959.

[16] G. Lorentz and R. Lorentz: *Solvability problems of bivariate interpolation I*, Constr. Approx., 2:153–169, 1986

[17] G. Migliorati: *Polynomial approximation by means of the random discrete L2 projection and application to inverse problems for PDEs with stochastic data*, Ph.D. thesis, Dipartimento di Matematica "Francesco Brioschi", Politecnico di Milano, Milano, Italy, and Centre de Mathématiques Appliquées, École Polytechnique, Palaiseau, France, 2013.

[18] G. Migliorati: *Multivariate Markov-type and Nikolskii-type inequalities for polynomials associated with downward closed multi-index sets*, J. Approx. Theory, 189:137–159, 2015.

[19] G. Migliorati and F. Nobile: *Analysis of discrete least squares on multivariate polynomial spaces with evaluations at low-discrepancy point sets*, J. Complexity, 31(4):517–542, 2015.

[20] G. Migliorati, F. Nobile, E. von Schwerin, and R. Tempone: *Approximation of Quantities of Interest in stochastic PDEs by the random discrete L^2 projection on polynomial spaces*, SIAM J. Sci. Comput., 35(3):A1440–A1460, 2013.

[21] G. Migliorati, F. Nobile, E. von Schwerin, and R. Tempone: *Analysis of discrete L^2 projection on polynomial spaces with random evaluations*, Found. Comput. Math., 14:419–456, 2014.

[22] G. Migliorati, F. Nobile, and R. Tempone: *Convergence estimates in probability and in expectation for discrete least squares with noisy evaluations at random points*, J. Multivar. Anal., 142:167–182, 2015.

[23] T. Poggio and S. Smale: *The mathematics of learning: Dealing with data*. Notices Amer. Math. Soc., 50:537–544, 2003.

[24] H. Rauhut and C. Schwab: *Compressive sensing Petrov-Galerkin approximation of high dimensional parametric operator equations*, Math. Comp., accepted, 2016.

[25] H. Rauhut and R. Ward: *Interpolation via weighted l1-minimization*, Appl. Comput. Harmon. Anal., 40:321–351, 2016.

[26] V.N. Temlyakov: *Approximation in Learning Theory*, Constr. Approx., 27:33–74, 2008.

[27] V. Vapnik: *Statistical learning theory*, John Wiley & Sons, 1998.

[28] T. Zhang: *Sparse recovery with orthogonal matching pursuit under RIP*, IEEE Trans. Inform. Theory, 57(9):6215–6221, 2011.